arXiv:1410.8175v1 [cs.SI] 29 Oct 2014

Randomized Rumor Spreading in Poorly Connected Small-World Networks*

Abbas Mehrabian†
Department of Combinatorics and Optimization,
University of Waterloo, Waterloo, Ontario, Canada
amehrabi@uwaterloo.ca

Ali Pourmiri
Max Planck Institute for Informatics, Saarbrücken, Germany
pourmiri@mpi-inf.mpg.de

October 31, 2014

Abstract

The Push-Pull protocol is a well-studied round-robin rumor spreading protocol defined as follows: initially a node knows a rumor and wants to spread it to all nodes in a network quickly. In each round, every informed node sends the rumor to a random neighbor, and every uninformed node contacts a random neighbor and gets the rumor from her if she knows it. We analyze the behavior of this protocol on random k-trees, a class of power law graphs, which are small-world and have large clustering coefficients, built as follows: initially we have a k-clique. In every step a new node is born, a random k-clique of the current graph is chosen, and the new node is joined to all nodes of the k-clique. When $k \geqslant 2$ is fixed, we show that if initially a random node is aware of the rumor, then with probability $1 - o(1)$ after $\mathcal{O}\left((\log n)^{1+2/k} \cdot \log\log n \cdot f(n)\right)$rounds the rumor propagates to $n - o(n)$ nodes, where n is the number of nodes and $f(n)$ is any slowly growing function. Since these graphs have polynomially small conductance, vertex expansion $\mathcal{O}(1/n)$ and

*A preliminary version of this paper has appeared in proceedings of the 28th International Symposium on Distributed Computing (DISC 2014), pages 346–360.

†Supported by the Vanier Canada Graduate Scholarships program. Part of this work was done while the first author was visiting Monash University, Australia.

constant treewidth, these results demonstrate that Push-Pull can be efficient even on poorly connected networks.

On the negative side, we prove that with probability $1 - o(1)$ the protocol needs at least $\Omega\left(n^{(k-1)/(k^2+k-1)}/f^2(n)\right)$ rounds to inform all nodes. This exponential dichotomy between time required for informing *almost all* and *all* nodes is striking. Our main contribution is to present, for the first time, a natural class of random graphs in which such a phenomenon can be observed. Our technique for proving the upper bound successfully carries over to a closely related class of graphs, the random k-Apollonian networks, for which we prove an upper bound of $\mathcal{O}\left((\log n)^{c_k} \cdot \log\log n \cdot f(n)\right)$ rounds for informing $n - o(n)$ nodes with probability $1 - o(1)$ when $k \geqslant 3$ is fixed. Here, $c_k = (k^2-3)/(k-1)^2 < 1 + 2/k$.

Keywords: randomized rumor spreading, push-pull protocol, random k-trees, random k-Apollonian networks, urn models.

1 Introduction

Randomized rumor spreading is an important primitive for information dissemination in networks and has numerous applications in network science, ranging from spreading information in the WWW and Twitter to spreading viruses and diffusion of ideas in human communities (see [5, 10, 11, 12, 19]). A well studied rumor spreading protocol is the *Push-Pull protocol*, introduced by Demers, Greene, Hauser, Irish, Larson, Shenker, Sturgis, Swinehart, and Terry [9]. Suppose that one node in a network is aware of a piece of information, the 'rumor.' The protocol proceeds in rounds. In each round, every *informed* node contacts a random neighbor and sends the rumor to it ('pushes' the rumor), and every *uninformed* nodes contacts a random neighbor and gets the rumor if the neighbor knows it ('pulls' the rumor). Note that this is a synchronous protocol, e.g. a node that receives a rumor in a certain round cannot send it on in the same round.

A point to point communication network can be modelled as an undirected graph: the nodes represent the processors and the links represent communication channels between them. Studying rumor spreading has several applications to distributed computing in such networks, of which we mention just two. The first is in broadcasting algorithms: a single processor wants to broadcast a piece of information to all other processors in the network (see [25] for a survey). There are at least three advantages to the Push-Pull protocol: it is simple (each node makes a simple local decision in each round; no knowledge of the global topology is needed; no state is maintained), scalable (the protocol is independent of the size of network: it does not grow more complex as the network grows) and robust (the protocol tolerates random node/link failures without the use of error recovery mechanisms, see [15]). A second application comes from the maintenance of

databases replicated at many sites, e.g., yellow pages, name servers, or server directories. There are updates injected at various nodes, and these updates must propagate to all nodes in the network. In each round, a processor communicates with a random neighbor and they share any new information, so that eventually all copies of the database converge to the same contents. See [9] for details. Other than the aforementioned applications, rumor spreading protocols have successfully been applied in various contexts such as resource discovery [24], distributed averaging [3], data aggregation [28], and the spread of computer viruses [2].

We only consider simple, undirected and connected graphs. For a graph G, let $\Delta(G)$ and $\mathrm{diam}(G)$ denote the maximum degree and the diameter of G, respectively, and let $\deg(v)$ denote the degree of a vertex v. Most studies in randomized rumor spreading focus on the *runtime* of this protocol, defined as the number of rounds taken until a rumor initiated by one vertex reaches all other vertices. It is clear that $\mathrm{diam}(G)/2$ is a lower bound for the runtime of the Push-Pull protocol. We say an event happens *with high probability (whp)* if its probability approaches 1 as n goes to infinity. Feige, Peleg, Raghavan and Upfal [15] showed that for an n-vertex G, whp the rumor reaches all vertices in $\mathcal{O}(\Delta(G) \cdot (\mathrm{diam}(G) + \log n))$ rounds. This protocol has been studied on many graph classes such as complete graphs [27], Erdős-Réyni random graphs [14, 15, 17, 35], random regular graphs [1, 18], and hypercube graphs [15]. For most of these classes it turns out that whp the runtime is $\mathcal{O}(\mathrm{diam}(G) + \log n)$, which does not depend on the maximum degree.

Randomized rumor spreading has recently been studied on real-world networks models. Doerr, Fouz, and Friedrich [10] proved an upper bound of $\mathcal{O}(\log n)$ for the runtime on preferential attachment graphs, and Fountoulakis, Panagiotou, and Sauerwald [19] proved the same upper bound (up to constant factors) for the runtime on the giant component of random graphs with given expected degrees (also known as the Chung-Lu model) with power law degree distribution.

The runtime is closely related to the *expansion profile* of the graph. Let $\Phi(G)$ and $\alpha(G)$ denote the conductance and the vertex expansion of a graph G, respectively. After a series of results by various scholars, Giakkoupis [22, 23] showed that for any n-vertex graph G, the runtime of the Push-Pull protocol is $\mathcal{O}\left(\min\{\Phi(G)^{-1} \cdot \log n, \alpha(G)^{-1} \cdot \log \Delta(G) \cdot \log n\}\right)$. It is known that whp preferential attachment graphs and random graphs with power law expected degrees have conductance $\Omega(1)$ (see [6, 32]). So it is not surprising that rumors spread fast on these graphs. Censor-Hillel, Haeupler, Kelner, and Maymounkov [4] presented a different rumor spreading protocol that whp distributes the rumor in $\mathcal{O}(\mathrm{diam}(G) + \mathrm{polylog}(n))$ rounds on any connected n-vertex graph, which seems particularly suitable for poorly connected graphs.

1.1 Our contribution

We study the Push-Pull protocol on random k-trees, a class of random graphs defined as follows.

Definition (Random k-tree process [20])**.** Let k be a positive integer. Build a sequence $G(0), G(1), \ldots$ of random graphs as follows. The graph $G(0)$ is just a clique on k vertices. For each $1 \leq t \leq n$, $G(t)$ is obtained from $G(t-1)$ as follows: a k-clique of $G(t-1)$ is chosen uniformly at random, a new vertex is born and is joined to all vertices of the chosen k-clique. The graph $G(n)$ is called a random k-tree on $n+k$ vertices.

We remark that this process is different from the random k-tree process defined by Cooper and Uehara [8] which was further studied in [7].

Sometimes it is convenient to view this as a 'random graph evolving in time.' In this interpretation, in every round $1, 2, \ldots$, a new vertex is born and is added to the evolving graph, and $G(t)$ denotes the graph at the end of round t. Observe that $G(t)$ has $k+t$ many vertices and $kt+1$ many k-cliques.

As in the preferential attachment scheme, the random k-tree process enjoys a 'the rich get richer' effect. Think of the number of k-cliques containing any vertex v as the 'wealth' of v (note that this quantity is linearly related to $\deg(v)$). Then, the probability that the new vertex attaches to v is proportional to the wealth of v, and if this happens, the wealth of v increases by $k-1$. On the other hand, random k-trees have much larger clustering coefficients than preferential attachment graphs, as all neighbors of each new vertex are joined to each other. It is well-known that real-world networks tend to have large clustering coefficients (see, e.g., [36, Table 1]).

Gao [20] showed that whp the degree sequence of $G(n)$ asymptotically follows a power law distribution with exponent $2 + \frac{1}{k-1}$. In Section 3 we show that whp the diameter of $G(n)$ is $\mathcal{O}(\log n)$, and its clustering coefficient is at least $1/2$, as opposed to preferential attachment graphs and random graphs with power law expected degrees, whose clustering coefficients are $o(1)$ whp. As per these properties, random k-trees may serve as more realistic models for real-world networks.

On the other hand, in Section 3 we prove that with high probability a random k-tree on $n+k$ vertices has conductance $\mathcal{O}\left(\log n \cdot n^{-1/k}\right)$ and vertex expansion $\mathcal{O}(k/n)$. Therefore we cannot resort to existing results linking the runtime to expansion properties to show rumors spread fast in these graphs. Another interesting structural property of a random k-tree is its *treewidth* (see [29] for a comprehensive survey). Gao [21] proved that many random graph models, including Erdős-Réyni random graphs with expected degree $\omega(\log n)$ and preferential attachment graphs with out-degree greater than 11, have treewidth $\Theta(n)$, whereas all random k-trees have treewidth k by construction. (According

to [21], not much is known about the treewidth of a preferential attachment graph with out-degree between 3 and 11.)

In conclusion, distinguishing features of random k-trees, such as high clustering coefficient, bad expansion (polynomially small conductance) and tree-like structure (small treewidth), inspired us to study randomized rumor spreading on this unexplored random environment. Our first main contribution is the following theorem.

Theorem 1. *Let $k \geqslant 2$ be constant and let $f(n) = o(\log\log n)$ be an arbitrary function going to infinity with n. If initially a random vertex of an $(n+k)$-vertex random k-tree knows a rumor, then with high probability after $\mathcal{O}\left((\log n)^{1+\frac{2}{k}} \cdot \log\log n \cdot f(n)^{\frac{3}{k}}\right)$ rounds of the* Push-Pull *protocol, $n - o(n)$ vertices will know the rumor.*

We give a high-level sketch of the proof of Theorem 1. Let $m = o(n)$ be a suitably chosen parameter, and note that $G(m)$ is a subgraph of $G = G(n)$. Consider the connected components of $G - G(m)$. Most vertices born later than round m have relatively small degree, so most these components have a small maximum degree (and logarithmic diameter) thus the rumor spreads quickly inside each of them. A vertex $v \in V(G(m))$ typically has a large degree, but this means there is a high chance that v has a neighbor x with small degree, which quickly receives the rumor from v and spreads it (or vice versa). We build an almost-spanning tree T of $G(m)$ with logarithmic height, such that for every edge uv of T, one of u and v have a small degree, or u and v have a common neighbor with a small degree. Either of these situations mean the rumor is exchanged quickly between u and v. This tree T then works as a 'highway system' to spread the rumor within vertices of $G(m)$ and from them to the components of $G - G(m)$.

The main novelty in this proof is how the almost-spanning tree is built and used (using small degree vertices for fast rumor transmission between high degree vertices has also been used in previous papers, e.g. [10, 19]). Our second main contribution is the following theorem, which gives a polynomial lower bound for the runtime.

Theorem 2. *Let $f(n) = o(\log\log n)$ be an arbitrary function going to infinity with n. Suppose that initially one vertex in the random k-tree, $G(n)$, knows the rumor. Then, with high probability the* Push-Pull *protocol needs at least $n^{(k-1)/(k^2+k-1)} f(n)^{-2}$ rounds to inform all vertices of $G(n)$.*

We give a high-level sketch of the proof of Theorem 2. A *barrier* in a graph is a subset D of edges of size $\mathcal{O}(1)$, whose deletion disconnects the graph. If both endpoints of every edge of a barrier D have large degrees, then the protocol needs a large time to pass the rumor through D. For proving Theorem 2, we prove a random k-tree has a barrier whp. The main novelty in this proof is introducing and using the notion of a barrier.

It is instructive to contrast Theorems 1 and 2. The former implies that if you want to inform almost all the vertices, then you just need to wait for a polylogarithmic number of rounds. The latter implies that, however, if you want to inform each and every vertex, then you have to wait for polynomially many rounds. This is a striking phenomenon and the main message of this paper is to present, for the first time, a natural class of random graphs in which this phenomenon can be observed. In fact, in applications such as viral marketing and voting, it is more appealing to inform 99 percent of the vertices very quickly instead of waiting a long time until everyone gets informed. For such applications, Theorem 1 implies that the Push-Pull protocol can be effective even on poorly connected graphs.

It is worth mentioning that bounds for the number of rounds to inform *almost all* vertices have already appeared in the literature, see for instance [11, 19]. In particular, for power-law Chung-Lu graphs with exponent in $(2, 3)$, it is shown in [19] that whp after $\mathcal{O}(\log\log n)$ rounds the rumor spreads in $n - o(n)$ vertices, but to inform *all* vertices of the giant component $\Theta(\log n)$ rounds are needed. This result also shows a great difference between the two cases, however in both cases the required time is quite small.

A closely related class of graphs is the class of *random k-Apollonian networks*, introduced by Zhang, Comellas, Fertin, and Rong [37]. Their construction is very similar to that of random k-trees, with just one difference: if a k-clique is chosen in a certain round, it will never be chosen again. It is known that whp random k-Apollonian networks exhibit a power law degree distribution and large clustering coefficient [34, 38] and have logarithmic diameter [7]. Our third main contribution is the following theorem.

Theorem 3. *Let $k \geqslant 3$ be constant and let $f(n) = o(\log\log n)$ be an arbitrary function going to infinity with n. Assume that initially a random vertex of an $(n+k)$-vertex random k-Apollonian network knows a rumor. Then, with high probability after*

$$\mathcal{O}\left((\log n)^{(k^2-3)/(k-1)^2} \cdot \log\log n \cdot f(n)^{2/k}\right)$$

rounds of the Push-Pull *protocol, at least $n - o(n)$ vertices will know the rumor.*

The proof of Theorem 3 is along the lines of that of Theorem 1, although there are several differences. Note that we have $(k^2 - 3)/(k - 1)^2 < 1 + 2/k$, so our upper bound for random k-Apollonian networks is slightly stronger than that for random k-trees.

Unfortunately, our technique for proving Theorem 2 does not extend to random k-Apollonian networks, although we believe that whp we need a polynomial number of rounds to inform all vertices in a random k-Apollonian network as well. We leave this as a conjecture.

For the rest of the paper, k is a constant greater than 1, and the asymptotics are for n going to infinity. Several times in our proofs we use urn models to analyze the vertices'

degrees and the number of vertices in certain parts of a random k-tree. We also use a result on the height of random recursive trees to conclude that random k-trees have logarithmic diameter. The connections with urn models are built in Section 2. In Section 3 we study basic properties of random k-trees, demonstrating their similarities with real-world graphs. Theorems 1, 2, and 3 are proved in Sections 4, 5, and 6, respectively.

2 Connections with urn theory

We will need some definitions and results from urn theory (see [26] for a general introduction). After reviewing these, we build some connections with random k-trees that will be used throughout.

Definition (Pólya-Eggenberger urn)**.** Start with W_0 white and B_0 black balls in an urn. In every step a ball is drawn from the urn uniformly at random, the ball is returned to the urn, and s balls of the same color are added to the urn. Let $\mathrm{Urn}(W_0, B_0, s, n)$ denote the number of white balls right after n draws.

Proposition 1. *Let $X = \mathrm{Urn}(a, b, k, n)$ and $w = a + b$. Then*

$$\mathbf{E}\left[X^2\right] = \left(a + \frac{a}{w}\, kn\right)^2 + \frac{abk^2n(kn+w)}{w^2(w+k)}$$

and for any $c \geq (a+b)/k$ we have

$$\mathbf{Pr}\left[X = a\right] \leq \left(\frac{c}{c+n}\right)^{a/k}.$$

Proof. The first statement follows from the following well known formulae for the expected value and the variance of X (see [31, Corollary 5.1.1] for instance):

$$\mathbf{E}\left[X\right] = a + \frac{a}{w}\, kn\,, \quad \mathbf{Var}\left[X\right] = \frac{abk^2n(kn+w)}{w^2(w+k)}\,.$$

For the second statement, we have

$$\begin{aligned}
\mathbf{Pr}\left[X=a\right] &= \frac{b}{a+b}\cdot\frac{b+k}{a+b+k}\cdot\cdots\cdot\frac{b+(n-1)k}{a+b+(n-1)k} \\
&= \prod_{i=0}^{n-1}\left(1-\frac{a}{a+b+ik}\right) \\
&\leqslant \prod_{i=0}^{n-1}\left(1-\frac{a}{ck+ik}\right) \\
&\leqslant \exp\left(-\sum_{i=0}^{n-1}\frac{a}{ck+ik}\right) \\
&= \left\{\exp\left(\sum_{i=0}^{n-1}\frac{1}{c+i}\right)\right\}^{-a/k} \leqslant \left\{\exp\left(\int_{x=c}^{c+n}\frac{\mathrm{d}x}{x}\right)\right\}^{-a/k} = \left(\frac{c}{c+n}\right)^{a/k}. \quad \square
\end{aligned}$$

Definition (Generalized Pólya-Eggenberger urn)**.** Let $\alpha, \beta, \gamma, \delta$ be nonnegative integers. We start with W_0 white and B_0 black balls in an urn. In every step a ball is drawn from the urn uniformly at random and returned to the urn. Additionally, if the ball is white, then δ white balls and γ black balls are returned to the urn; otherwise, i.e. if the ball is black, then β white balls and α black balls are returned to the urn. Let $\mathrm{Urn}\left(W_0, B_0, \begin{bmatrix}\alpha & \beta\\ \gamma & \delta\end{bmatrix}, n\right)$ denote the number of white balls right after n draws.

Note that Pólya-Eggenberger urns correspond to the matrix $\begin{bmatrix}s & 0\\ 0 & s\end{bmatrix}$. The following proposition follows from known results.

Proposition 2. *Let* $X = \mathrm{Urn}\left(W_0, B_0, \begin{bmatrix}\alpha & 0\\ \gamma & \delta\end{bmatrix}, n\right)$ *and let* r *be a positive integer. If* $\gamma, \delta > 0$*,* $\alpha = \gamma + \delta$*, and* $r\delta \geqslant \alpha$*, then we have*

$$\mathbf{E}\left[X^r\right] \leqslant \left(\frac{\alpha n}{W_0+B_0}\right)^{r\delta/\alpha}\prod_{i=0}^{r-1}\left(W_0+i\delta\right) + \mathcal{O}\left(n^{(r-1)\delta/\alpha}\right).$$

Proof. By [16, Proposition 15] we have

$$\mathbf{E}\left[X^r\right] = n^{r\delta/\alpha}\delta^r\frac{\Gamma(W_0/\delta+r)\Gamma((W_0+B_0)/\alpha)}{\Gamma(W_0/\delta)\Gamma((W_0+B_0+r\delta)/\alpha)} + \mathcal{O}\left(n^{(r-1)\delta/\alpha}\right),$$

Note that

$$\frac{\Gamma(W_0/\delta+r)}{\Gamma(W_0/\delta)} = \prod_{i=0}^{r-1}\left(i+W_0/\delta\right).$$

Finally, the inequality

$$\frac{\Gamma((W_0+B_0+r\delta)/\alpha)}{\Gamma((W_0+B_0)/\alpha)} \geqslant ((W_0+B_0)/\alpha)^{r\delta/\alpha}$$

follows from $r\delta \geqslant \alpha$ and the following inequality (see, e.g., [30, inequality (2.2)])

$$\frac{\Gamma(x+1)}{\Gamma(x+s)} \geqslant x^{1-s} \qquad \forall x>0, s\in[0,1]\,. \quad \square$$

Proposition 3. *Suppose that in $G(j)$ vertex x has $A>0$ neighbors, and is contained in B many k-cliques. Conditional on this, the degree of x in $G(n+j)$ is distributed as*

$$A+\left(\mathrm{Urn}\left(B, kj+1-B, \begin{bmatrix} k & 0 \\ 1 & k-1 \end{bmatrix}, n\right) - B\right) \Big/ (k-1)\,.$$

Proof. We claim that the total number of k-cliques containing x in $G(n+j)$ is distributed as $\mathrm{Urn}\left(B, kj+1-B, \begin{bmatrix} k & 0 \\ 1 & k-1 \end{bmatrix}, n\right)$. At the end of round j, there are B many k-cliques containing x, and $kj+1-B$ many k-cliques not containing x. In each subsequent round $j+1,\dots,j+n$, a random k-clique is chosen and k new k-cliques are created. If the chosen k-clique contains x, then $k-1$ new k-cliques containing x are created, and 1 new k-clique not containing x is created. Otherwise, i.e. if the chosen k-clique does not contain x, then no new k-cliques containing x is created, and k new k-cliques not containing x are created, and the claim follows.

Hence the number of k-cliques that are created in rounds $j+1,\dots,j+n$ and contain x is distributed as $\mathrm{Urn}\left(B, kj+1-B, \begin{bmatrix} k & 0 \\ 1 & k-1 \end{bmatrix}, n\right) - B$, and the proof follows by noting that every new neighbor of x creates $k-1$ new k-cliques containing x. $\square$

Combining Propositions 2 and 3 we obtain the following lemma.

Lemma 4. *Let $1 \leqslant j \leqslant n$ and let q be a positive integer. Let x denote the vertex born in round j. Conditional on any $G(j)$, the probability that x has degree greater than $k+q(n/j)^{(k-1)/k}$ in $G(n)$ is $\mathcal{O}\left(q\sqrt{q}\exp(-q)\right)$.*

Proof. Let $X = \mathrm{Urn}\left(k, kj-k+1, \begin{bmatrix} k & 0 \\ 1 & k-1 \end{bmatrix}, n-j\right)$. By Proposition 3, $\deg(x)$ is distributed as $k+(X-k)/(k-1)$. By Proposition 2,

$$\mathbf{E}\left[X^q\right] \leqslant (1+o(1))\left(\frac{k(n-j)}{kj+1}\right)^{\frac{q(k-1)}{k}} \prod_{i=0}^{q-1}(k+i(k-1)) \leqslant \left(\frac{n}{j}\right)^{\frac{q(k-1)}{k}} (k-1)^q(q+1)!\,.$$

Thus,

$$\begin{aligned}\mathbf{Pr}\left[\deg(x) > k + q(n/j)^{(k-1)/k}\right] &= \mathbf{Pr}\left[X - k > q(k-1)(n/j)^{(k-1)/k}\right] \\ &\leqslant \frac{\mathbf{E}\left[X^q\right]}{\left(q(k-1)(n/j)^{(k-1)/k}\right)^q} \\ &\leqslant (q+1)!q^{-q} = \mathcal{O}\left(q\sqrt{q}\exp(-q)\right) . \end{aligned}$$ □

3 Basic properties of random k-trees

In this section we prove that random k-trees exhibit two important properties observed in real-world networks: low diameter and large clustering coefficient. We also prove that random k-trees do not expand well, confirming our claim in the introduction that random k-trees are poorly connected graphs and thus existing techniques do not apply. Let $G(0), G(1), \ldots$ be defined as in Definition 1.1.

Definition. The *clustering coefficient* of a graph G, written $cc(G)$, is defined as

$$cc(G) = \frac{1}{|V(G)|} \sum_{u \in V(G)} \frac{|\langle N(u) \rangle|}{\binom{\deg(u)}{2}} ,$$

where $|\langle N(u) \rangle|$ denotes the number of edges xy such that both x and y are neighbors of u.

Proposition 5. *For every positive integer n, deterministically, the clustering coefficient of $G(n)$ is at least 1/2.*

Proof. Let u be a vertex of $G = G(n)$. It is not hard to check that $|\langle N(u) \rangle| = (k-1)(\deg(u) - k/2)$, and since $\deg(u) \geqslant k$ we get

$$\frac{|\langle N(u) \rangle|}{\binom{\deg(u)}{2}} \geqslant \frac{k}{\deg(u)} .$$

Using the Cauchy-Schwarz inequality we get

$$cc(G) \geqslant \frac{1}{|V(G)|} \sum_{u \in V(G)} \frac{k}{\deg(u)} \geqslant \frac{k}{n+k} \cdot \frac{(n+k)^2}{2|E(G)|} \geqslant \frac{1}{2} .$$ □

For proving random k-trees have logarithmic diameter we will need a known result about random d-ary recursive trees.

Definition (Random d-ary recursive tree)**.** Let d be a positive integer. Build a sequence $T(0), T(1), \ldots$ of rooted random trees as follows. The tree $T(0)$ has just one vertex, the root. For each $1 \leq t \leq n$, $T(t)$ is obtained from $T(t-1)$ as follows: a leaf of $T(t-1)$ is chosen uniformly at random and gives birth to d new children. The tree $T(n)$ is called a random d-ary recursive tree on $dn+1$ vertices.

Theorem 4 ([13],Theorem 6.47)**.** *Let α be the unique solution in (d, ∞) of*

$$\alpha(d-1)\log\left(\frac{de}{\alpha(d-1)}\right) = 1\,.$$

Let H_n denote the height of a random d-ary recursive tree on $dn+1$ vertices. There exists a constant $c > 0$ such that for any η,

$$\mathbf{Pr}\left[H_n > \alpha\log n + \eta\right] = \mathcal{O}(e^{-c\eta})\,.$$

The following proposition implies that with high probability the diameter of $G(n)$ is $\mathcal{O}(\log n)$.

Proposition 6. *Whp $G(n)$ has the following property: let $u_h u_{h-1}\cdots u_0$ be an arbitrary path such that u_i is born later than u_{i-1} for all i; then $h = \mathcal{O}(\log n)$.*

Proof. We inductively define a notion of *draft* for vertices and k-cliques of $G(n)$. The draft of the vertices of $G(0)$ as well as the k-clique they form equals 0. The draft of every k-clique equals the maximum draft of its vertices. Whenever a new vertex is born and is joined to a k-clique, the draft of the vertex equals the draft of the k-clique plus one. It is easy to see that if $xy \in E(G(n))$ and x is born later than y, then $\mathrm{draft}(x) \geqslant \mathrm{draft}(y)+1$. In particular, if x is a vertex of $G(n)$ and there is a path $x = x_h, x_{h-1}, \ldots, x_1, x_0$ in $G(n)$ such that x_j is born later than x_{j-1} for each j, then $\mathrm{draft}(x) \geqslant h$. Hence we just need to show that with high probability the draft of each k-clique of $G(n)$ is $\mathcal{O}(\log n)$.

We define an auxiliary tree whose nodes correspond to the k-cliques of $G(n)$, in such a way that the depth of each node in this tree equals the draft of its corresponding k-clique. Start with a single node corresponding to $G(0)$. Whenever a new vertex x is born and is joined to a k-clique C, k new k-cliques are created. In the auxiliary tree, add these to the set of children of C. So, the auxiliary tree evolves as follows: in every round a node is chosen uniformly at random and gives birth to k new children. Hence, the height of the auxiliary tree after n rounds is stochastically smaller than that of a random k-ary recursive tree on $1+kn$ nodes, whose height is $\mathcal{O}(\log n)$ whp by Theorem 4. □

Definition. The *vertex expansion* of a graph G (also known as the *vertex isoperimetric number* of G), written $\alpha(G)$, is defined as

$$\alpha(G) = \min\left\{\frac{|\partial S|}{|S|} : S \subseteq V(G), 0 < |S| \leq |V(G)|/2\right\}\,,$$

where ∂S denotes the set of vertices in $V(G) \setminus S$ that have a neighbor in S.

Definition. The *conductance* of a graph G, written $\Phi(G)$, is defined as

$$\Phi(G) = \min \left\{ \frac{e(S, V(G) \setminus S)}{\texttt{vol}(S)} : S \subseteq V(G), 0 < \texttt{vol}(S) \leqslant \texttt{vol}(V(G))/2 \right\} ,$$

where $e(S, V(G) \setminus S)$ denotes the number of edges between S and $V(G) \setminus S$, and $\texttt{vol}(S) = \sum_{u \in S} \deg(u)$ for every $S \subseteq V(G)$.

Proposition 7. *Deterministically $G(n)$ has vertex expansion $\mathcal{O}(k/n)$, and whp its conductance is $\mathcal{O}\left(\log n \cdot n^{-1/k}\right)$.*

Proof. Let $G = G(n)$. Since G has treewidth k, by [29, Lemma 5.3.1] there exists a partition (A, B, C) of $V(G)$ such that

1. $|C| = k + 1$,
2. $(n-1)/3 \leqslant |A| \leqslant 2(n-1)/3$ and $(n-1)/3 \leqslant |B| \leqslant 2(n-1)/3$, and
3. there is no edge between A and B.

At least one of A and B, say A, has size less than $(n+k)/2$. Then

$$\alpha(G) \leqslant \frac{|\partial A|}{|A|} \leqslant \frac{k+1}{(n-1)/3} = \mathcal{O}(k/n) .$$

At least one of A and B, say B, has volume less than $\texttt{vol}(G)/2$. Then since all vertices in G have degrees at least k,

$$\Phi(G) \leqslant \frac{e(B, A \cup C)}{\texttt{vol}(B)} \leqslant \frac{e(B, C)}{k|B|} \leqslant \frac{(k+1)\Delta(G)}{k(n-1)/3} = \mathcal{O}(\Delta(G)/n) .$$

Hence to prove $\Phi(G) = \mathcal{O}\left(\log n \cdot n^{-1/k}\right)$ it suffices to show that with high probability we have

$$\Delta(G) \leqslant k + (2 \log n) n^{1-1/k} .$$

Let $q = \lfloor 2 \log n \rfloor$ and let x be a vertex born in one of the rounds $1, 2, \ldots, n$. By Lemma 4,

$$\mathbf{Pr}\left[\deg(x) > k + q n^{1-1/k}\right] = \mathcal{O}(q\sqrt{q}\exp(-q)) = o(1/n) .$$

An argument similar to the proof of Lemma 4 shows that the probability that a vertex in $G(0)$ has degree greater than $k + q n^{1-1/k}$ is $o(1/n)$ as well. A union bound over all vertices shows that with high probability we have $\Delta(G) \leqslant k + (2 \log n) n^{1-1/k}$, as required. $\square$

4 Proof of Theorem 1

Once we have the following lemma, our problem reduces to proving a structural result for random k-trees.

Lemma 8. *Let χ and τ be fixed positive integers. Let G be an n-vertex graph and let $\Sigma \subseteq V(G)$ with $|\Sigma| = n - o(n)$ be such that for every pair of vertices $u, v \in \Sigma$ there exists a (u,v)-path $uu_1u_2\cdots u_{l-1}v$ such that $l \leqslant \chi$ and for every $0 \leqslant i \leqslant l-1$ we have $\min\{\deg(u_i), \deg(u_{i+1})\} \leqslant \tau$ (where we define $u_0 = u$ and $u_l = v$). If a random vertex in G knows a rumor, then whp after $6\tau(\chi + \log n)$ rounds of the* Push-Pull *protocol, at least $n - o(n)$ vertices will know the rumor.*

Proof. The proof is along the lines of that of [15, Theorem 2.2]. We show that given any $u, v \in \Sigma$, if u knows the rumor then with probability at least $1 - o\left(n^{-2}\right)$ after $6\tau(\chi + \log n)$ rounds v will know the rumor. The lemma follows by using the union bound and noting that a random vertex lies in Σ with high probability. Consider the (u,v)-path $uu_1u_2\cdots u_{l-1}v$ promised by the hypothesis. We bound from below the probability that the rumor is passed through this path.

For every $0 \leq i \leq l-1$, the number of rounds taken for the rumor to pass from u_i to u_{i+1} is a geometric random variable with success probability at least $1/\tau$ (if $\deg(u_i) \leqslant \tau$, this is the number of rounds needed for u_i to push the rumor along the edge, and if $\deg(u_{i+1}) \leqslant \tau$, this is the number of rounds needed for u_{i+1} to pull the rumor along the edge). The random variables corresponding to distinct edges are mutually independent. Hence the probability that the rumor is not passed in $6\tau(\chi + \log n)$ rounds is at most the probability that the number of heads in a sequence of $6\tau(\chi + \log n)$ independent biased coin flips, each having probability $1/\tau$ of being heads, is less than l. Let X denote the number of heads in such a sequence. Then using the Chernoff bound (see, e.g., [33, Theorem 4.2]) and noting that $\mathbf{E}[X] = 6(\chi + \log n)$ we get

$$\mathbf{Pr}[X < l] \leqslant \mathbf{Pr}[X \leq \mathbf{E}[X]/6] \leq \exp(-(5/6)^2 \mathbf{E}[X]/2) \leq \exp(-(5/6)^2(6\log n)/2),$$

which is $o\left(n^{-2}\right)$, as required. □

Let $f(n) = o(\log\log n)$ be an arbitrary function going to infinity with n, and let

$$m = \left\lceil \frac{n}{f(n)^{3/(k-1)}(\log n)^{2/(k-1)}} \right\rceil .$$

Also let $q = \lceil 4\log\log n \rceil$ and let

$$\tau = 2k + q(n/m)^{1-1/k} . \tag{1}$$

By Proposition 6, whp a random k-tree on $n+k$ vertices has diameter $\mathcal{O}(\log n)$. Theorem 1 thus follows from Lemma 8 and the following structural result, which we prove in the rest of this section.

Lemma 9. *Let G be an $(n+k)$-vertex random k-tree. Whp there exists $\Sigma \subseteq V(G)$ satisfying the conditions of Lemma 8 with τ defined in (1) and $\chi = \mathcal{O}(\log n + \operatorname{diam}(G))$.*

For the rest of this section, G is an $(n+k)$-vertex random k-tree. Recall from Definition 1.1 that $G = G(n)$, where $G(0), G(1), \dots$ is the random k-tree process. Consider the graph $G(m)$, which has $k+m$ vertices and $mk+1$ many k-cliques. For an edge e of $G(m)$, let $N(e)$ denote the number of k-cliques of $G(m)$ containing e. We define a spanning forest F of $G(m)$ as follows: for every $1 \leqslant t \leqslant m$, if the vertex x born in round t is joined to the k-clique C, then in F, x is joined to a vertex $u \in V(C)$ such that

$$N(xu) = \max_{v \in V(C)} N(xv)\,.$$

Note that F has k trees and the k vertices of $G(0)$ lie in distinct trees. Think of these trees as rooted at these vertices. The tree obtained from F by merging these k vertices is the 'highway system' described in the sketch of the proof of Theorem 1. Informally speaking, the proof has three parts: first, we show that this tree has a small height (Lemma 10); second, we show that each edge in this tree quickly exchanges the rumor with a reasonably large probability (Lemma 12); and finally we show that almost all vertices in $G - G(m)$ have quick access to and from F (Lemma 13).

Let LOG denote the event 'each tree in F has height $\mathcal{O}(\log n)$.' The following lemma is an immediate corollary of Proposition 6.

Lemma 10. *With high probability* LOG *happens.*

We prove Lemma 9 conditional on the event LOG. In fact, we prove it for any $G(m)$ that satisfies LOG. Let G_1 be an arbitrary instance of $G(m)$ that satisfies LOG. So, G_1 and F are fixed in the following, and all randomness refers to rounds $m+1, \dots, n$. The following deterministic lemma will be used in the proof of Lemma 12.

Lemma 11. *Assume that $xy \in E(F)$ and x is born later than y. If the degree of x in G_1 is greater than $2k-2$, then $N(xy) \geqslant (k^2-k)/2$.*

Proof. Assume that x is joined to $u_1, \dots, u_k$ when it is born, and that $v_1, v_2, \dots, v_{k-1}, \dots$ are the neighbors of x that are born later than x, in the order of birth. Let Ψ denote the number of pairs (u_j, C), where $1 \leqslant j \leqslant k$, and C is a k-clique in G_1 containing the edge xu_j. Consider the round in which vertex x is born and is joined to $u_1, \dots, u_k$. For

every $j \in \{1, \dots, k\}$, the vertex u_j is contained in $k-1$ new k-cliques, so in this round Ψ increases by $k(k-1)$. For each $i \in \{1, \dots, k-1\}$, consider the round in which vertex v_i is born. This vertex is joined to x and $k-1$ neighbors of x. At this round x has neighbor set $\{u_1, \dots, u_k, v_1, \dots, v_{i-1}\}$. Thus at least $k-i$ of the u_j's are joined to v_i in this round. Each vertex u_j that is joined to v_i in this round is contained in $k-2$ new k-cliques that contain x as well, so in this round Ψ increases by at least $(k-i)(k-2)$. Consequently, we have

$$\Psi \geqslant k(k-1) + \sum_{i=1}^{k-1} (k-i)(k-2) = k^2(k-1)/2\,.$$

By the pigeonhole principle, there exists some $\ell \in \{1, \dots, k\}$ such that the edge xu_ℓ is contained in at least $(k^2-k)/2$ many k-cliques, and this completes the proof. □

A vertex of G is called *modern* if it is born later than the end of round m, and is called *traditional* otherwise. In other words, vertices of G_1 are traditional and vertices of $G - G_1$ are modern. We say edge $uv \in E(G)$ is *fast* if at least one of the following is true: $\deg(u) \leqslant \tau$, or $\deg(v) \leqslant \tau$, or u and v have a common neighbor w with $\deg(w) \leqslant \tau$. For an edge $uv \in E(F)$, let $p_S(uv)$ denote the probability that uv is not fast, and let p_S denote the maximum of p_S over all edges of F.

Lemma 12. *We have $p_S = o(1/(f(n)\log n))$.*

Proof. Let $xy \in E(F)$ be arbitrary. By symmetry we may assume that x is born later than y. By Lemma 11, at least one of the following is true: vertex x has less than $2k-1$ neighbors in G_1, or $N(xy) \geqslant (k^2-k)/2$. So we may consider two cases.

- Case 1: vertex x has less than $2k-1$ neighbors in G_1. In this case vertex x lies in at most k^2-2k+2 many k-cliques of G_1. Assume that x has A neighbors in G_1 and lies in B many k-cliques in G_1. Let

$$X = \mathrm{Urn}\left(B, km+1-B, \begin{bmatrix} k & 0 \\ 1 & k-1 \end{bmatrix}, n-m\right)\,.$$

  Then by Proposition 3 the degree of x is distributed as $A + (X-B)/(k-1)$. By Proposition 2,

$$\begin{aligned}
\mathbf{E}\left[X^q\right] &\leqslant (1+o(1))\left(\frac{k(n-m)}{km+1}\right)^{\frac{q(k-1)}{k}} \prod_{i=0}^{q-1}(B+i(k-1)) \\
&\leqslant (1+o(1))\left(\frac{n}{m}\right)^{\frac{q(k-1)}{k}} (k-1)^q \prod_{i=0}^{q-1}(k+i) \\
&\leqslant (k-1)^q (k+q)!\left(\frac{n}{m}\right)^{\frac{q(k-1)}{k}}\,,
\end{aligned}$$

where we have used $B \leqslant k(k-1)$ for the second inequality. Therefore,

$$\mathbf{Pr}\left[\deg(x) > 2k + q(n/m)^{\frac{k-1}{k}}\right] \leqslant \mathbf{Pr}\left[X \geqslant (k-1)q(n/m)^{\frac{k-1}{k}}\right]$$

$$\leqslant \frac{\mathbf{E}\left[X^q\right]}{(k-1)^q q^q (n/m)^{\frac{q(k-1)}{k}}} = \mathcal{O}\left(\frac{(k+q)^{k+q}\sqrt{q}}{q^q \exp(k+q)}\right) = o\left(\frac{1}{f(n)\log n}\right) .$$

- Case 2: $N(xy) \geqslant (k^2-k)/2$. In this case we bound from below the probability that there exists a modern vertex w that is adjacent to x and y and has degree at most τ. We first bound from above the probability that x and y have no modern common neighbors. For this to happen, none of the k-cliques containing x and y must be chosen in rounds $m+1, \ldots, n$. This probability equals $\mathbf{Pr}\left[\mathrm{Urn}(N(xy), mk+1-N(xy), k, n-m) = N(xy)\right]$. Since $N(xy) \geqslant (k^2-k)/2$, by Proposition 1 we have

  $$\mathbf{Pr}\left[\mathrm{Urn}(N(xy), mk+1-N(xy), k, n-m) = N(xy)\right] \leqslant \left(\frac{m+1}{n+1}\right)^{\frac{k-1}{2}} ,$$

  which is $o\left(1/(f(n)\log n)\right)$.

  Now, assume that x and y have a modern common neighbor w. If there are multiple such vertices, choose the one that is born first. Since w appears later than round m, by Lemma 4,

  $$\mathbf{Pr}\left[\deg(w) > k + q(n/m)^{(k-1)/k}\right] = \mathcal{O}\left(q\sqrt{q}\exp(-q)\right) = o\left(\frac{1}{f(n)\log n}\right) . \quad \square$$

Enumerate the k-cliques of G_1 as $C_1, \ldots, C_{mk+1}$. Then choose $r_1 \in C_1, \ldots, r_{mk+1} \in C_{mk+1}$ arbitrarily, and call them the *representative vertices*. Starting from G_1, when modern vertices are born in rounds $m+1, \ldots, n$ until G is formed, every clique C_i 'grows' to a random k-tree with a random number of vertices, which is a subgraph of G. Enumerate these subgraphs as $H_1, \ldots, H_{mk+1}$, and call them the *pieces*. More formally, $H_1, \ldots, H_{mk+1}$ are induced subgraphs of G such that a vertex v is in $V(H_j)$ if and only if every path connecting v to a traditional vertex intersects $V(C_j)$. In particular, $V(C_j) \subseteq V(H_j)$ for all $j \in \{1, \ldots, mk+1\}$. Note that the H_j's may intersect, as a traditional vertex may lie in more than one C_j, however every modern vertex lies in a unique piece.

A traditional vertex is called *nice* if it is connected to some vertex in $G(0)$ via a path of fast edges. Since F has height $\mathcal{O}(\log n)$ and each edge of F is fast with probability at least $1-p_S$, the probability that a given traditional vertex is not nice is $\mathcal{O}(p_S \log n)$ by the union bound. A piece H_j is called *nice* if all its modern vertices have degrees at most τ, and the vertex r_j is nice. A modern vertex is called *nice* if it lies in a nice piece. A vertex/piece is called *bad* if it is not nice.

Lemma 13. *The expected number of bad vertices is $o(n)$.*

Proof. The total number of traditional vertices is $k + m = o(n)$ so we may just ignore them in the calculations below. Let $\eta = nf(n)/m = o(\log^3 n)$. Say piece H_j is *sparse* if $|V(H_j)| \leqslant \eta + k$. We first bound the expected number of modern vertices in non-sparse pieces. Observe that the number of modern vertices in a given piece is distributed as $X = (\mathrm{Urn}(1, km, k, n - m) - 1)/k$. Using Proposition 1 we get $\mathbf{E}[X^2] \leqslant 2kn^2/m^2$. By the second moment method, for every $t > 0$ we have

$$\mathbf{Pr}[X \geqslant t] \leqslant \frac{\mathbf{E}[X^2]}{t^2} \leqslant \frac{2kn^2}{m^2t^2}\,.$$

The expected number of modern vertices in non-sparse pieces is thus at most

$$\begin{aligned}(km+1)\sum_{i=0}^{\infty}(2^{i+1}\eta)\,\mathbf{Pr}\left[2^i\eta < X \leqslant 2^{i+1}\eta\right] &\leqslant \sum_{i=0}^{\infty}(2^{i+1}\eta)(km+1)\frac{2kn^2}{m^2\eta^2 2^{2i}} \\ &\leqslant \mathcal{O}\left(\frac{n^2}{m\eta}\right)\sum_{i=0}^{\infty}2^{-i} = \mathcal{O}\left(\frac{n^2}{m\eta}\right)\,,\end{aligned}$$

which is $o(n)$.

We now bound the expected number of modern vertices in sparse bad pieces. For bounding this from above we find an upper bound for the expected number of bad pieces, and multiply by η. A piece H_j can be bad in two ways:

(1) the representative vertex r_j is bad: the probability of this is $\mathcal{O}(p_S \log n)$. Therefore, the expected number of pieces that are bad due to this reason is $\mathcal{O}(mkp_S \log n)$, which is $o(n/\eta)$ by Lemma 12.

(2) there exists a modern vertex in H_j with degree greater than τ: the probability that a given modern vertex has degree greater than τ is $\mathcal{O}\left(q\sqrt{q}\exp(-q)\right)$ by Lemma 4. So the average number of modern vertices with degree greater than τ is $\mathcal{O}\left(nq\sqrt{q}\exp(-q)\right)$. Since every modern vertex lies in a unique piece, the expected number of pieces that are bad because of this reason is bounded by $\mathcal{O}\left(nq\sqrt{q}\exp(-q)\right) = o(n/\log^3 n)$.

So the expected number of bad pieces is $o(n/\eta + n/\log^3 n)$, and the expected number of modern vertices in sparse bad pieces is $o(n + \eta n/\log^3 n) = o(n)$. □

We now prove Lemma 9, which concludes the proof of Theorem 1.

Proof of Lemma 9. Let Σ denote the set of nice modern vertices. By Lemma 13 and using Markov's inequality, we have $|\Sigma| = n - o(n)$ whp. Let $\{a_1, \ldots, a_k\}$ denote the vertex set of $G(0)$. Using an argument similar to the proof of Lemma 12, it can be proved that

given $1 \leqslant i < j \leqslant k$, the probability that edge $a_i a_j$ is not fast is $o(1)$. Since the total number of such edges is a constant, whp all such edges are fast. Let u and v be nice modern vertices, and let r_u and r_v be the representative vertices of the pieces containing them, respectively. Since the piece containing u is nice, there exists a (u, r_u)-path whose vertices except possibly r_u all have degrees at most τ. The length of this path is at most $\mathrm{diam}(G)$. Since r_u is nice, for some $1 \leqslant i \leqslant n$ there exists an (r_u, a_i)-path in F consisting of fast edges. Appending these paths gives a (u, a_i)-path with length at most $\mathrm{diam}(G) + \mathcal{O}(\log n)$ such that for every pair of consecutive vertices in this path, one of them has degree at most τ. Similarly, for some $1 \leqslant j \leqslant n$ there exists a (v, a_j)-path of length $\mathcal{O}(\log n + \mathrm{diam}(G))$, such that one of every pair of consecutive vertices in this path has degree at most τ. Since the edge $a_i a_j$ is fast whp, we can build a (u, v)-path of length $\mathcal{O}(\log n + \mathrm{diam}(G))$ of the type required by Lemma 8, and this completes the proof. □

5 Proof of Theorem 2

Definition (s-barrier)**.** A pair $\{C_1, C_2\}$ of disjoint k-cliques in a connected graph is an *s-barrier* if (i) the set of edges between C_1 and C_2 is a cut-set, i.e. deleting them disconnects the graph, and (ii) the degree of each vertex in $V(C_1) \cup V(C_2)$ is at least s.

Observe that if G has an s-barrier, then for any starting vertex, whp the Push-Pull protocol needs at least $\Omega(s)$ rounds to inform all vertices.

Lemma 14. *The graph $G(n)$ has an $\Omega(n^{1-1/k})$-barrier with probability $\Omega(n^{1/k-k})$.*

Proof. Let $u_1, \ldots, u_k$ be the vertices of $G(0)$, and let $v_1, \ldots, v_k$ be the vertices of $G(k) - G(0)$ in the order of appearance. We define two events. Event A is that for every $1 \leqslant i \leqslant k$, when v_i appears, it attaches to $v_1, v_2, \ldots, v_{i-1}, u_i, u_{i+1}, \ldots, u_k$; and for each $1 \leqslant i, j \leqslant k$, u_i and v_j have no common neighbor in $G(n) - G(k)$. Event B is that all vertices of $G(k)$ have degree $\Omega(n^{(k-1)/k})$ in $G(n)$. Note that if A and B both happen, then the pair $\{u_1 u_2 \ldots u_k, v_1 v_2 \ldots v_k\}$ is an $\Omega(n^{(k-1)/k})$-barrier in $G(n)$. To prove the lemma we will show that $\mathbf{Pr}[A] = \Omega(n^{1/k-k})$ and $\mathbf{Pr}[B|A] = \Omega(1)$.

For A to happen, first, the vertices $v_1, \ldots, v_k$ must choose the specified k-cliques, which happens with constant probability. Moreover, the vertices appearing after round k must not choose any of the $k^2 - 1$ many k-cliques that contain both u_i's and v_j's. Since

$1 - y \geqslant e^{-y-y^2}$ for every $y \in [0, 1/4]$,

$$\begin{aligned}
\mathbf{Pr}\left[A\right] &= \Omega(\mathbf{Pr}\left[\mathrm{Urn}(k^2-1, 2, k, n-k) = k^2-1\right]) \\
&= \Omega\left(\prod_{i=0}^{n-k-1}\left(\frac{2+ik}{k^2+1+ik}\right)\right) \\
&\geqslant \Omega\left(\prod_{i=0}^{4k-1}\left(\frac{2+ik}{k^2+1+ik}\right)\prod_{i=4k}^{n-k-1}\left(1-\frac{k^2-1}{ik}\right)\right) \\
&\geqslant \Omega\left(\exp\left(-\sum_{i=4k}^{n-k-1}\left\{\frac{k^2-1}{ik}+\left(\frac{k^2-1}{ik}\right)^2\right\}\right)\right)
\end{aligned}$$

which is $\Omega(n^{1/k-k})$ since

$$\sum_{i=4k}^{n-k-1}\frac{k^2-1}{ik} \leqslant (k-1/k)\log n + \mathcal{O}(1) \text{ and } \sum_{i=4k}^{n-k-1}\left(\frac{k^2-1}{ik}\right)^2 = \mathcal{O}(1)\,.$$

Conditional on A and using an argument similar to that in the proof of Proposition 3, the degree of each of the vertices $u_1, \ldots, u_k, v_1, \ldots, v_k$ in $G(n)$ is at least $k + (\mathrm{Urn}(1, 1, \begin{bmatrix} k & 0 \\ 1 & k-1 \end{bmatrix}, n-k) - 1)/(k-1)$. By [16, Proposition 16], there exists $\delta > 0$ such that

$$\mathbf{Pr}\left[\mathrm{Urn}(1, 1, \begin{bmatrix} k & 0 \\ 1 & k-1 \end{bmatrix}, n-k) < \delta n^{(k-1)/k}\right] < 1/(2k+1)\,.$$

By the union bound, the probability that all vertices $u_1, \ldots, u_k, v_1, \ldots, v_k$ have degrees at least $\delta n^{(k-1)/k}/(k-1)$ is at least $1/(2k+1)$, hence $\mathbf{Pr}\left[B|A\right] \geqslant 1/(2k+1) = \Omega(1)$. □

Let $f(n) = o(\log\log n)$ be any function going to infinity with n, and also let $m = \left\lceil f(n)n^{1-k/(k^2+k-1)}\right\rceil$. (Note that the value of m is different from that in Section 4, although its role is somewhat similar.) Consider the random k-tree process up to round m. Enumerate the k-cliques of $G(m)$ as $C_1, \ldots, C_{mk+1}$. Starting from $G(m)$, when new vertices are born in rounds $m+1, \ldots, n$ until $G = G(n)$ is formed, every clique C_i 'grows' to a random k-tree with a random number of vertices, which is a subgraph of G. Enumerate these subgraphs as $H_1, \ldots, H_{mk+1}$, and call them the *pieces*. We say a piece is *moderate* if its number of vertices is between $n/(mf(n))$ and $nf(n)/m$. Note that the number of vertices in a piece has expected value $\Theta(n/m)$. The following lemma is proved by showing this random variable does not deviate much from its expected value.

Lemma 15. *With high probability, there are $o(m)$ non-moderate pieces.*

Proof. We prove the first piece, H_1, is moderate whp. By symmetry, this would imply that the average number of non-moderate pieces is $o(m)$. By Markov's inequality, this gives that whp there are $o(m)$ non-moderate pieces. Let X denote the number of vertices of H_1. Note that X is distributed as $k + \mathrm{Urn}(1, km, k, n-m)$; so its expected value is $k + \frac{n-m}{1+km} = \Theta(n/m)$. By Markov's inequality, $\mathbf{Pr}[X > nf(n)/m] = o(1)$.

For bounding $\mathbf{Pr}[X < n/(mf(n))]$, we use an alternative way to define the random variable $\mathrm{Urn}(1, km, k, n-m)$ (see [26, page 181]): assume Z is a beta random variable with parameters $1/k$ and m. Then $X-k$, which has the same distribution as $\mathrm{Urn}(1, km, k, n-m)$, is distributed as a binomial random variable with parameters $n-m$ and Z. Note that

$$\begin{aligned}\mathbf{Pr}[Z < 3/(mf(n))] &= \frac{\Gamma(m+1/k)}{\Gamma(m)\Gamma(1/k)} \int_0^{3/(mf(n))} x^{1/k-1}(1-x)^{m-1}dx \\ &< \frac{m^{1/k}}{\Gamma(1/k)} \int_0^{3/(mf(n))} x^{1/k-1}dx = \frac{3^{1/k}k}{\Gamma(1/k)f(n)^{1/k}} = o(1)\,,\end{aligned}$$

where we have used the fact $\Gamma(m+1/k) < \Gamma(m)m^{1/k}$ which follows from [30, inequality (2.2)]. On the other hand, the Chernoff bound (see, e.g., [33, Theorem 4.2]) gives

$$\begin{aligned}\mathbf{Pr}[X < n/(mf(n)) | Z \geqslant 3/(mf(n))] &\leqslant \mathbf{Pr}[\mathrm{Bin}(n-m, 3/(mf(n))) < n/(mf(n))] \\ &< \exp(-3(n-m)/(8mf(n))) = o(1)\,,\end{aligned}$$

thus $\mathbf{Pr}[X < n/(mf(n))] = o(1)$. □

Proof of Theorem 2. Consider an alternative way to generate $G(n)$ from $G(m)$: first, we determine how many vertices each piece has, and then we expose the structure of the pieces. Let Y denote the number of moderate pieces. By Lemma 15 we have $Y = \Omega(m)$ whp. We prove the theorem conditional on $Y = y$, where $y = \Omega(m)$ is otherwise arbitrary. Note that after the sizes of the pieces are exposed, what happens inside each piece in rounds $m+1, \ldots, n$ is mutually independent from other pieces. Let H be a moderate piece with n_1 vertices. By Lemma 14, the probability that H has an $\Omega(n_1^{1-1/k})$-barrier is $\Omega(n_1^{1/k-k})$. Since $n/(mf(n)) \leqslant n_1 \leqslant nf(n)/m$, the probability that H has a $\Omega((n/(mf(n))^{1-1/k})$-barrier is $\Omega((nf(n)/m)^{1/k-k})$. Since there are $y = \Omega(m)$ moderate pieces in total, the probability that no moderate piece has an $\Omega\left((n/(mf(n)))^{1-1/k}\right)$-barrier is at most

$$(1 - \Omega((nf(n)/m)^{1/k-k}))^y \leqslant \exp(-\Omega(f(n))) = o(1)\,,$$

so whp there exists an $\Omega\left(n^{(k-1)/(k^2+k-1)}f(n)^{-2}\right)$-barrier in $G(n)$, as required. □

6 Proof of Theorem 3

In this section we analyze the Push-Pull protocol on random k-Apollonian networks. Since these networks are a sub-family of random k-trees, we can reuse the proof techniques in Section 4 to bound the time needed to inform almost all vertices. First, we formally define the random k-Apollonian process.

Definition (Random k-Apollonian process)**.** Let k be a positive integer. Build a sequence $A(0)$, $A(1)$, ... of random graphs as follows. The graph $A(0)$ is just a clique on k vertices. This k-clique is marked as *active*. For each $1 \leq t \leq n$, $A(t)$ is obtained from $A(t-1)$ as follows: an active k-clique of $A(t-1)$ is chosen uniformly at random, a new vertex is born and is joined to all vertices of the chosen k-clique. The chosen k-clique is marked as *non-active*, and all the new k-cliques are marked as active in $A(t)$. The graph $A(n)$ is called a *random k-Apollonian network (k-RAN)* on $n+k$ vertices.

We first prove a counterpart of Lemma 4 for k-RANs.

Lemma 16. *Let $1 \leqslant j \leqslant n$ and let q be a positive integer. Let x denote the vertex born in round j. Conditional on any $A(j)$, the probability that x has degree greater than $k + q(n/j)^{(k-2)/(k-1)}$ in $A(n)$ is $\mathcal{O}\left(q\sqrt{q}\exp(-q)\right)$.*

Proof. Let $X = \mathrm{Urn}\left(k, (k-1)(j-1), \begin{bmatrix} k-1 & 0 \\ 1 & k-2 \end{bmatrix}, n-j\right)$. An argument similar to the proof of Proposition 3 gives that the degree of x in $A(n)$ is distributed as $k + (X-k)/(k-2)$. By Proposition 2,

$$\mathbf{E}\left[X^q\right] \leqslant (1+o(1))\left(\frac{(k-1)(n-j)}{(k-1)j+1}\right)^{\frac{q(k-2)}{k-1}} \prod_{i=0}^{q-1}(k+i(k-2)) \leqslant \left(\frac{n}{j}\right)^{\frac{q(k-2)}{k-1}} (k-2)^q(q+1)!.$$

Thus,

$$\begin{aligned}\mathbf{Pr}\left[\deg(x) > k + q(n/j)^{(k-2)/(k-1)}\right] &= \mathbf{Pr}\left[X - k > (k-2)q(n/j)^{(k-2)/(k-1)}\right] \\ &\leqslant \frac{\mathbf{E}\left[X^q\right]}{\left((k-2)q(n/j)^{(k-2)/(k-1)}\right)^q} \\ &\leqslant (q+1)!q^{-q} = \mathcal{O}\left(q\sqrt{q}\exp(-q)\right) .\end{aligned}$$

□

Fix $k > 2$ and let $f(n) = o(\log\log n)$ be an arbitrary function going to infinity with n, and let

$$m = \left\lceil \frac{n}{(\log n)^{2/(k-1)} f(n)^{(2k-2)/(k^2-2k)}} \right\rceil .$$

Finally, let $q = \lceil 4 \log\log n \rceil$ and let

$$\tau = 2k + q(n/m)^{(k-2)/(k-1)} \,. \tag{2}$$

An argument similar to the proof of Lemma 18 gives that whp a k-RAN on $n+k$ vertices has diameter $\mathcal{O}(\log n)$. Theorem 3 thus follows from Lemma 8 and the following structural result, which we prove in the rest of this section.

Lemma 17. *Let A be an $(n+k)$-vertex k-RAN. Whp there exists $\Sigma \subseteq V(A)$ satisfying the conditions of Lemma 8 with τ defined in (2) and $\chi = \mathcal{O}(\log n + \mathrm{diam}(A))$.*

The proof of Lemma 17 is along the lines of that of Lemma 9. For the rest of this section, $A = A(n)$ is an $(n+k)$-vertex k-RAN. Consider the graph $A(m)$, which has $k+m$ vertices and $m(k-1)+1$ active k-cliques. For any edge e of $A(m)$, let $N^*(e)$ denote the number of active k-cliques of $A(m)$ containing e. Note that, since $k > 2$, for each edge e, the number of active k-cliques containing e does not decrease as the k-RAN evolves. We define a spanning forest F of $A(m)$ as follows: at round 0, F has k isolated vertices, i.e. the vertices of $A(0)$; then for every $1 \leqslant t \leqslant m$, if the vertex x born in round t is joined to the k-clique C, then in F, x is joined to a vertex $u \in V(C)$ such that

$$N^*(xu) = \max_{v \in V(C)} N^*(xv) \,.$$

Note that F has k trees and the k vertices of $A(0)$ lie in distinct trees. Let LOG denote the event 'each tree in F has height $\mathcal{O}(\log n)$.'

Lemma 18. *With high probability* LOG *happens.*

Proof. We prove that whp every path $u_h u_{h-1} \cdots u_0$ in $A(n)$ such that u_i is born later than u_{i-1} for all i, has length $\mathcal{O}(\log n)$. The proof is very similar to that of Proposition 6, the only difference being that the built auxiliary tree is indeed a random k-ary recursive tree, whose height is whp $\mathcal{O}(\log n)$ by Theorem 4. □

We prove Lemma 17 conditional on the event LOG. In fact, we prove it for any $A(m)$ that satisfies LOG. Let A_1 be an arbitrary instance of $A(m)$ that satisfies LOG. So, A_1 and F are fixed in the following, and all randomness refers to rounds $m+1, \ldots, n$. The following deterministic lemma will be used in the proof of Lemma 20.

Lemma 19. *Assume that $xy \in E(F)$ and x is born later than y. If the degree of x in A_1 is at least $2k-1$, then $N^*(xy) \geqslant (k-1)^2/2$.*

Proof. Assume that x is joined to $u_1, \dots, u_k$ when it is born, and that $v_1, v_2, \dots, v_{k-1}, \dots$ are the neighbors of x that are born later than x, in the order of birth. Let Ψ denote the number of pairs (u_j, C), where C is an active k-clique in A_1 with $xu_j \subseteq E(C)$. Consider the round in which vertex x is born and is joined to $u_1, \dots, u_k$. For every $j \in \{1, \dots, k\}$, the edge xu_j is contained in $k-1$ new active k-cliques, so in this round Ψ increases by $k(k-1)$. For each $i \in \{1, \dots, k-1\}$, consider the round in which vertex v_i is born. At least $k-i$ of the u_j's are joined to v_i in this round. Each vertex u_j that is joined to v_i in this round is contained in $k-2$ new k-cliques that contain x, and one k-clique containing u_j is deactivated. Hence in this round Ψ increases by at least $(k-i)(k-3)$. Consequently, right after v_{k-1} is born, we have

$$\Psi \geqslant k(k-1) + \sum_{i=1}^{k-1} (k-i)(k-3) = (k-1)^2 k/2 \,.$$

By the pigeonhole principle, there exists some $\ell \in \{1, \dots, k\}$ such that the edge xu_ℓ is contained in at least $(k-1)^2/2$ active k-cliques, and this completes the proof, as the number of active k-cliques containing xu_ℓ will not decrease later. □

A vertex of A is called *modern* if it is born later than the end of round m, and is called *traditional* otherwise. In other words, vertices of A_1 are traditional and vertices of $A - A_1$ are modern. We say edge $uv \in E(A)$ is *fast* if at least one of the following is true: $\deg(u) \leqslant \tau$, or $\deg(v) \leqslant \tau$, or u and v have a common neighbor w with $\deg(w) \leqslant \tau$. For an edge $uv \in E(F)$, let $p_S(uv)$ denote the probability that uv is not fast, and let p_S denote the maximum of p_S over all edges of F.

Lemma 20. *We have* $p_S = o(1/(f(n) \log n))$.

Proof. The proof is similar to that of Lemma 12. Let $xy \in E(F)$ be arbitrary. By symmetry we may assume that x is born later than y. By Lemma 19, at least one of the following is true: vertex x has less than $2k-1$ neighbors in A_1, or $N^*(xy) \geqslant (k-1)^2/2$. So we may consider two cases.

- Case 1: vertex x has less than $2k-1$ neighbors in A_1. In this case x lies in at most $k + (k-2)^2$ many active k-cliques of A_1. Suppose that x has D neighbors in A_1 and lies in B many active k-cliques in A_1. Let

$$X = \mathrm{Urn}\left(B, (k-1)m + 1 - B, \begin{bmatrix} k-1 & 0 \\ 1 & k-2 \end{bmatrix}, n-m\right) .$$

  Then by an argument similar to the proof of Proposition 3, the degree of x is

distributed as $D + (X - B)/(k-2)$. By Proposition 2,

$$\begin{aligned}\mathbf{E}[X^q] &\leqslant (1+o(1))\left(\frac{(k-1)(n-m)}{(k-1)m+1}\right)^{\frac{q(k-2)}{k-1}} \prod_{i=0}^{q-1}(B+i(k-2)) \\ &\leqslant \mathcal{O}\left(\left(\frac{n}{m}\right)^{\frac{q(k-2)}{k-1}}(k-2)^q(k+q)!\right),\end{aligned}$$

where we have used $B \leqslant k(k-2)$. Therefore,

$$\begin{aligned}\mathbf{Pr}\left[\deg(x) > 2k + q(n/m)^{\frac{k-2}{k-1}}\right] &\leqslant \mathbf{Pr}\left[X \geqslant (k-2)q(n/m)^{\frac{k-2}{k-1}}\right] \\ &\leqslant \frac{\mathbf{E}[X^q]}{(k-2)^q q^q (n/m)^{\frac{q(k-2)}{k-1}}} = \mathcal{O}\left(\frac{(k+q)!}{q^q}\right) = o\left(\frac{1}{f(n)\log n}\right).\end{aligned}$$

- Case 2: $N^*(xy) \geqslant (k-1)^2/2$. In this case we bound from below the probability that there exists a modern vertex w that is adjacent to x and y and has degree at most τ. We first bound from above the probability that x and y have no modern common neighbors. For this to happen, none of the k-cliques containing x and y must be chosen in rounds $m+1, \ldots, n$. This probability equals

$$p := \mathbf{Pr}\left[\mathrm{Urn}(N^*(xy), m(k-1)+1-N^*(xy), k-1, n-m) = N^*(xy)\right].$$

Since $N^*(xy) \geqslant (k-1)^2/2$, by Proposition 1 we have

$$p \leqslant \left(\frac{m+1}{n}\right)^{(k-1)/2} = o\left(\frac{1}{f(n)\log n}\right).$$

Now, assume that x and y have a modern common neighbor w. If there are multiple such vertices, choose the one that is born first. Since w appears later than round m, by Lemma 16,

$$\mathbf{Pr}\left[\deg(w) > k + q(n/m)^{(k-2)/(k-1)}\right] = \mathcal{O}\left(q\sqrt{q}\exp(-q)\right) = o\left(\frac{1}{f(n)\log n}\right). \quad \square$$

Enumerate the k-cliques of A_1 as $C_1, C_2, \ldots$, and $C_{m(k-1)+1}$. Then choose $r_1 \in C_1, \ldots, r_{m(k-1)+1} \in C_{m(k-1)+1}$ arbitrarily, and call them the *representative vertices*. Starting from A_1, when modern vertices are born in rounds $m+1, \ldots, n$ until A is formed, every clique C_i 'grows' to a k-RAN with a random number of vertices, which is a subgraph of A. Enumerate these subgraphs as $H_1, \ldots, H_{m(k-1)+1}$, and call them the *pieces*. More formally, $H_1, \ldots, H_{m(k-1)+1}$ are induced subgraphs of A such that a vertex v is in $V(H_j)$ if and only if every path connecting v to a traditional vertex intersects $V(C_j)$.

A traditional vertex is called *nice* if it is connected to some vertex in $A(0)$ via a path of fast edges. Since F has height $\mathcal{O}(\log n)$ and each edge of F is fast with probability at least $1 - p_S$, the probability that a given traditional vertex is not nice is $\mathcal{O}(p_S \log n)$ by the union bound. A piece H_j is called *nice* if all its modern vertices have degrees at most τ, and the vertex r_j is nice. A modern vertex is called *nice* if it lies in a nice piece. A vertex/piece is called *bad* if it is not nice.

Lemma 21. *The expected number of bad vertices is $o(n)$.*

Proof. The proof is very similar to that of Lemma 13, except we use Lemmas 16 and 20 instead of Lemmas 4 and 12, respectively. □

The proof of Lemma 17 is exactly the same as that of Lemma 9, except we use Lemmas 20 and 21 instead of Lemmas 12 and 13, respectively. This concludes the proof of Theorem 3.

References

[1] P. Berenbrink, R. Elsässer, and T. Friedetzky. Efficient randomised broadcasting in random regular networks with applications in peer-to-peer systems. In *Proc. 27th Symp. Principles of Distributed Computing (PODC)*, pages 155–164, 2008.

[2] N. Berger, C. Borgs, J.T. Chayes, and A. Saberi. On the spread of viruses on the Internet. In *Proc. 16th Symp. Discrete Algorithms (SODA)*, pages 301–310, 2005.

[3] Stephen Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6):2508–2530, 2006.

[4] Keren Censor-Hillel, Bernhard Haeupler, Jonathan A. Kelner, and Petar Maymounkov. Global computation in a poorly connected world: fast rumor spreading with no dependence on conductance. In *44th Symp. Theory of Computing (STOC)*, pages 961–970, 2012.

[5] Flavio Chierichetti, Silvio Lattanzi, and Alessandro Panconesi. Rumor spreading in social networks. In *Proc. 36th Intl. Coll. Automata, Languages and Programming (ICALP)*, pages 375–386, 2009.

[6] Fan R. K. Chung, Linyuan Lu, and Van H. Vu. The spectra of random graphs with given expected degrees. *Internet Mathematics*, 1(3):257–275, 2003.

[7] Colin Cooper and Alan Frieze. The height of random k-trees and related branching processes. *arXiv*, 1309.4342v2 [math.CO], 2013.

[8] Colin Cooper and Ryuhei Uehara. Scale free properties of random k-trees. *Mathematics in Computer Science*, 3(4):489–496, 2010.

[9] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. In *Proc. 6th Symp. Principles of Distributed Computing (PODC)*, pages 1–12, 1987.

[10] Benjamin Doerr, Mahmoud Fouz, and Tobias Friedrich. Social networks spread rumors in sublogarithmic time. In *Proc. 43th Symp. Theory of Computing (STOC)*, pages 21–30, 2011.

[11] Benjamin Doerr, Mahmoud Fouz, and Tobias Friedrich. Asynchronous rumor spreading in preferential attachment graphs. In *Proc. 13th Scandinavian Workshop Algorithm Theory (SWAT)*, pages 307–315, 2012.

[12] Benjamin Doerr, Mahmoud Fouz, and Tobias Friedrich. Why rumors spread so quickly in social networks. *Commun. ACM*, 55(6):70–75, 2012.

[13] Michael Drmota. *Random trees: An interplay between combinatorics and probability*. SpringerWienNewYork, Vienna, 2009.

[14] Robert Elsässer. On the communication complexity of randomized broadcasting in random-like graphs. In *Proceedings of the 18th ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA'06, pages 148–157, 2006.

[15] Uriel Feige, David Peleg, Prabhakar Raghavan, and Eli Upfal. Randomized broadcast in networks. *Random Struct. Algorithms*, 1(4):447–460, 1990.

[16] Philippe Flajolet, Philippe Dumas, and Vincent Puyhaubert. Some exactly solvable models of urn process theory. In *4th Colloquium on Mathematics and Computer Science Algorithms, Trees, Combinatorics and Probabilities*, Discrete Math. Theor. Comput. Sci. Proc., AG, pages 59–118. Assoc. Discrete Math. Theor. Comput. Sci., Nancy, 2006.

[17] Nikolaos Fountoulakis, Anna Huber, and Konstantinos Panagiotou. Reliable broadcasting in random networks and the effect of density. In *Proc. 29th IEEE Conf. Computer Communications (INFOCOM)*, pages 2552–2560, 2010.

[18] Nikolaos Fountoulakis and Konstantinos Panagiotou. Rumor spreading on random regular graphs and expanders. In *Proc. 14th Intl. Workshop on Randomization and Comput. (RANDOM)*, pages 560–573, 2010.

[19] Nikolaos Fountoulakis, Konstantinos Panagiotou, and Thomas Sauerwald. Ultra-fast rumor spreading in social networks. In *23th Symp. Discrete Algorithms (SODA)*, pages 1642–1660, 2012.

[20] Yong Gao. The degree distribution of random k-trees. *Theor. Comput. Sci.*, 410(8-10):688–695, 2009.

[21] Yong Gao. Treewidth of Erdős-Rényi random graphs, random intersection graphs, and scale-free random graphs. *Discrete Applied Mathematics*, 160(4-5):566–578, 2012.

[22] G. Giakkoupis. Tight bounds for rumor spreading with vertex expansion. In *Proc. 25th Symp. Discrete Algorithms (SODA)*, pages 801–815, 2014.

[23] George Giakkoupis. Tight bounds for rumor spreading in graphs of a given conductance. In *28th International Symposium on Theoretical Aspects of Computer Science (STACS 2011)*, volume 9, pages 57–68, 2011.

[24] Mor Harchol-Balter, Frank Thomson Leighton, and Daniel Lewin. Resource discovery in distributed networks. In *Proc. 18th Symp. Principles of Distributed Computing (PODC)*, pages 229–237, 1999.

[25] Sandra M. Hedetniemi, Stephen T. Hedetniemi, and Arthur L. Liestman. A survey of gossiping and broadcasting in communication networks. *Networks*, 18(4):319–349, 1988.

[26] Norman L. Johnson and Samuel Kotz. *Urn models and their application: An approach to modern discrete probability theory, Wiley Series in Probability and Mathematical Statistics*. John Wiley & Sons, New York-London-Sydney, 1977.

[27] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized Rumor Spreading. In *41st Symp. Foundations of Computer Science (FOCS)*, pages 565–574, 2000.

[28] David Kempe, Alin Dobra, and Johannes Gehrke. Gossip-based computation of aggregate information. In *44th Symp. Foundations of Computer Science (FOCS)*, pages 482–491, 2003.

[29] T. Kloks. *Treewidth: Computations and Approximations*. Springer-Verlag, 1994.

[30] A. Laforgia. Further inequalities for the gamma function. *Mathematics of Computation*, 42(166):597–600, 1984.

[31] Hosam M. Mahmoud. Pólya urn models and connections to random trees: A review. *Journal of the Iranian Statistical Society*, 2(1):53–114, 2003.

[32] Milena Mihail, Christos H. Papadimitriou, and Amin Saberi. On certain connectivity properties of the internet topology. In *Proc. 44th Symp. Foundations of Computer Science (FOCS)*, pages 28–35, 2003.

[33] R. Motwani and P. Raghavan. *Randomized algorithms*. Cambridge University Press, Cambridge, 1995.

[34] M. Mungan. Comment on "apollonian networks: Simultaneously scale-free, small world, Euclidean, space filling, and with matching graphs". *Phys. Rev. Lett.*, 106:029802, Jan 2011.

[35] K. Panagiotou, X. Pérez-Giménez, T. Sauerwald, and H. Sun. Randomized rumour spreading: The effect of the network topology. *Combinatorics, Probability and Computing*, FirstView:1–23, 5 2014.

[36] D. J. Watts and D. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393:440–442, 1998.

[37] Z. Zhang, F. Comellas, G. Fertin, and L. Rong. High-dimensional Apollonian networks. *J. Phys. A*, 39(8):1811–1818, 2006.

[38] T. Zhou, G. Yan, and B.-H. Wang. Maximal planar networks with large clustering coefficient and power-law degree distribution. *Phys. Rev. E*, 71:046141, Apr 2005.